arXiv:1704.01871v1 [stat.ML] 6 Apr 2017

Massive Data Clustering in Moderate Dimensions from the Dual Spaces of Observation and Attribute Data Clouds

Fionn Murtagh
Email: fmurtagh@acm.org

April 7, 2017

Abstract

Cluster analysis of very high dimensional data can benefit from the properties of such high dimensionality. Informally expressed, in this work, our focus is on the analogous situation when the dimensionality is moderate to small, relative to a massively sized set of observations. Mathematically expressed, these are the dual spaces of observations and attributes. The point cloud of observations is in attribute space, and the point cloud of attributes is in observation space. In this paper, we begin by summarizing various perspectives related to methodologies that are used in multivariate analytics. We draw on these to establish an efficient clustering processing pipeline, both partitioning and hierarchical clustering.

1 Introduction

From the next section, up to section 16, we summarize methodological perspectives. Alternatively expressed, we briefly describe parts of analytics processing paths, or alternative parts of such analytics processing pipelines.

- In section 2 there is a short introduction to the notation used.
- Section 3 covers the traditional approach of facilitating and expediting analysis through the forming of a reduced dimensionality accurate representative, and best fit, of the given data's dimensionality. Random projections can approximate well this processing.
- In section 4, there is (i) data piling that occurs in very high dimensions, (ii) random projections, but here shown to be very related and therefore this leads to interest in their aggregate or consensus.
- Since convergence or a sufficiency (e.g. of the aggregate of random projections) may be relevant, section 5 just notes how iterative convergence can

determine the principal eigenvector and eigenvalue, and hence the cloud's principal axis.

- The centring of the point clouds is one perspective on a trivial first eigenvector. Another perspective can be how Correspondence Analysis provides a mapping of the data relating to all that differentiates the data from a null χ^2 independence statistic. This is summarized in section 6.

- Section 7 describes how a hierarchic clustering, or tree structuring of interrelationships in the data, or a point set in an ultrametric, or tree, topology, can be easily mapped onto a one-dimensional alignment or ranking or curve.

- While fidelity of representation is one issue with low dimensional mapping, in section 8, more at issue is the one-dimensional alignment or seriation that could be usable for a hierarchical clustering, or, otherwise expressed, an ultrametric or tree topology embedding.

- The properties of the Baire metric, and simultaneously, ultrametric, are summarized in section 9.

- Section 10 notes how, here, partitions are assembled into a hierarchy, whereas traditionally, partitions are determined from a hierarchy.

- Further refining the hierarchical clustering that is associated with endowing the data with the Baire metric is considered in section 11. Benefits can be: storage; and interpretability (of the clustering analytics being carried out).

- Many aspects before section 12 relate to very high dimensional spaces. Given the dual space relationship, in this section interest is shifted to massive sized data sets.

- One way to exploit the dual space relationships is to infer in the high dimensional cloud from processing carried out in the lower dimensional cloud. Section 13.

- In section 14, at issue is use of data piling in one of the clouds, but not the other.

- How this all relates to seriation, the unidimensional mapping that might serve for reading off from it, a hierarchical clustering, i.e. an ultrametric or tree topology embedding, is summarized in section 15.

- Section 16 discusses the mapping of seriation into hierarchical clustering. Section 17 starts to deal with practical implementation. Section 18 pursues a case study in detail.

2 Introduction: Dual Spaces of Clouds of Points, Dimensionalities

Clouds of points are at issue here, where the observation points cloud is crossed by the attribute points clouds. Let I be the index set of the observation cloud, $\mathcal{N}(I)$, let J be the index set of the attribute cloud, $\mathcal{N}(J)$, and we have the observation set $x_I \subset \mathbb{R}^{|J|}$ and the attribute set $x_J \subset \mathbb{R}^{|I|}$. So the observation vector, x_i is a $|J|$-valued vector. Also the attribute vector, x_j is a $|I|$-valued vector.

Key processing elements in the description to follow are the normalizing or scaling carried out, on our data, and the mappings or projections, that are carried out. Both following on from the processing stages, and also arising from the initially obtained or sourced data, this both implies, and is directly related to the data distributional properties. A further important distinctive property, in the description to follow, is the case when cardinality is very large. This is, to begin with in the following sections, when the ambient dimensionality of our observations is of very large dimensionality, i.e. $|J|$ is very large, and $|J| \ll |I|$. Then we consider when the observation set is very large, i.e. $|I|$ is huge, and $|I| \ll |J|$.

Previous applications include chemical data (high dimensional) [14]; astronomy data (univariate, the aim being linear computational time hierarchical clustering), [3]; document similarity data [15]; protein clustering in computational biology, and enterprise information retrieval, [2].

3 Low Dimensional Mapping from High Dimensions through Random Projection

First, we note this aim of accurately mapping high dimensional data into a low dimensional space. By "accurate" is means preserving inter-point distances.

Random mapping of a high dimensional cloud of points into a lower dimensional subspace, with very little distortion of (dis)similarity, is described in [9]. This aim is to consider points $g, h \in \mathbb{R}^{\nu}$, i.e. ν-dimensional, that will be mapped, respectively, onto $x, y \in \mathbb{R}^{m}$, by means of a random linear transformation. Each mapped point in $\mathbb{R}^{m}$ is transformed to be of unit norm. Sought is to have $m \ll \nu$. Kaski [9] cites Hecht-Nielsen as follows: the number of *almost orthogonal directions* in a coordinate system, that is determined at random in a high dimensional space, is very much greater than the number of orthogonal directions. What is examined is the effect of the linear transformation, from the original space $\mathbb{R}^{\nu}$ to the new, randomly transformed space, $\mathbb{R}^{m}$. Distances in the original space, $\mathbb{R}^{\nu}$ are close to being fully preserved in the transformed space, $\mathbb{R}^{m}$. A theoretical statement of this is the Johnson-Lindenstrauss Lemma [8].

That linear transformation, $\mathbb{R}^{\nu} \longrightarrow \mathbb{R}^{m}$ uses vectors of zero mean, unit standard deviation Gaussian. Such a transformation, in [10], is termed the case

of conventional random projections.

4 Highly Correlated Random Projections of High Dimensional Point Clouds

Data concentration, also termed data piling, is a known characteristic of very high dimensional spaces. By the central limit theorem, and by the concentration (data piling) effect of high dimensions [7, 19], we have as dimension $\rightarrow \infty$: pairwise distances become equidistant; orientation tends to be uniformly distributed. Random projections are here carried out using uniformly, $[0, 1]$, distributed random axis coordinates. The projections are rescaled to the (closed/open) unit interval, $\geq 0, < 1$. We find high correlation between any pair of our random projections. See [15]. Unlike the conventional random mapping into a low dimensional space, as described in section 3, here we are using a uniform distributed mapping rather than a Gaussian distributed mapping. Also different from the conventional random mapping, projections are rescaled to the 0,1 interval, compared to normalizing to unit norm.

In the conventional random mapping, the aim is approximating distance preservation (termed the Parseval relation) in the low dimensional mapping. Rather than this, with our highly correlated random mappings, we seek instead to exploit and draw benefit from this mapping of clouds that are subject to data piling.

In order to do this, we form the consensus random projection. This is constituted by the mean of each point's random projected values.

5 Iterative Determining of the Principal Eigenvector

In [17], there is consideration of the iterative determination of the principal eigenvector. Such an algorithm is easily and straightforwardly implemented. (However it is computationally less interesting and relevant, compared to eigenreduction through diagonalization.) Also considered in [17] is power iteration clustering, basing the clustering on the principal axis projections.

There is discussion in [1] of such iterative eigenreduction, with an unbounded observations set, hence indicated as an infinite number of observations or input data table rows.

6 Correspondence Analysis Decomposition: Trivial Eigenvectors

In the eigen-decomposition, we have the following expressing, for $N = |I|$, and for factors, F, G.

$$f_{ij} = f_i f_j \left(1 + \sum_{\alpha=1,\ldots N} \lambda_\alpha^{-\frac{1}{2}} F_{\alpha(i)} G_{\alpha(j)} \right) \quad (1)$$

Here the first term on the right gives rise in practice to trivial eigenvalues, equal to 1, and to associated trivial eigenvectors. If we have, through data piling, a very large number of factors, in the summation, then it comes about that there is an approximation of the data, f_{ij} by $f_i f_j$. This, therefore, is the product of the marginal distributions. This, in turn, is given by a constant times the row and the column sums.

It may be just noted, that this amounts to a 0-valued χ^2 statistic, that fully supports expected values, defined from the row and column margins, specifying the observed data. With data piling, or data compactification, we can, and will, arrive at such an outcome.

7 Hierarchical Clustering and Seriation

A dendrogram's terminal nodes are a permutation of the object set that is clustered. Let us denote that permutation of the object (or observation) set, I, as $\pi(I)$. While of course the pairwise distances defined on I determines the hierarchical clustering, it is also very clear that that many variants on these pairwise distances would also determine an identical hierarchical clustering. Thus, one important characterizing aspect of a hierarchical clustering is $\pi(I)$, and its interpoint adjacent distances that lead to the hierarchical clustering. In [4], there is the foundation for hierarchical clustering requiring a seriation of what is to be clustered.

It is this seriation therefore, that we use, based on the consensus, that is, the highly correlated, set of random axis projections resulting from data piling, or compactification.

8 Two Distinct Objectives of Low Dimensionality Embedding

Compared to section 3 where the aim was the approximating of pairwise distance invariance in a low dimensional space embedding, here we bypass this objective in our analytics. From one perspective, such distance invariance is of benefit if we wish to carry out best match or nearest neighbour processing, with the approximation of every pairwise distance in the original space, and in the embedded space.

Arising out of how (see [4]) what is hierarchically clustered can be perfectly scaled in one dimension, we want to use this knowledge to proceed directly to our clusters. These clusters, being hierarchically structured, are determined on most, if not all, resolution scales. That is due to the taking full account of cluster embedding, or cluster inclusion, properties.

Here, therefore, our overriding objective is not so much fidelity in a new data space relative to our original data. Rather, we are concerned with both effective and efficient data scaling.

One form of realising such an objective is as follows. Our given data clouds are endowed with an appropriate metric, such as for positive counts or other real measurements, the χ^2 metric. That is, the data clouds are endowed with the Euclidean metric. Distance invariance will hold between (i) χ^2 distance in input space, and (ii) Euclidean distance in the output, factor space. That output space, by design, is also referred to as comprising: factors, principal components, principal axes, eigenvectors.

9 Baire or Longest Common Prefix Distance

Consider the consensus, viz. aggregated set of projections. Given two values, $x_{ik}, x_{i'k}$ for digits $k = 1, 2, ..., K$, the Baire distance of base B is defined as B^s where $s = \text{argmax}_k \, x_{ik} = x_{i'k}$.

In view of potential use of different number systems, B should be the number base that is in use, e.g. $B = 10$ for decimal and real numbers.

See [3] for the effective distance such that cluster members can be directly read off the data, once we specify each cluster, as a bin, specified by the set of shared, common prefix values.

The Baire, or longest common prefix distance, is both a metric and an ultrametric.

A property of the use of this metric, and ultrametric, is that we can specify any cluster, using the common prefix. Then we can read off the cluster members from our data. That is, a linear scan of the data is carried out to determine all members of a given cluster, specified by what is to be the common prefix. Thus, by endowing one's data with this metric and ultrametric, there is direct reading of the cluster members.

10 Practical Strategies for Hierarchical Clustering

Traditionally, one use of agglomerative hierarchical clustering has been to determine, for n observations, the extraction of one of the $n - 1$ partitions that are defined by the hierarchy. Motivation for that is lack of prior knowledge of number of clusters in a targeted partition.

Given the direct reading of clusters that are hierarchically structured, for very large datasets, it is beneficial to also extract one or more partitions. These partitions are on a range of resolution scales.

Taking our hierarchy as a regular B-way tree, for example with $B = 10$, then the first level of the hierarchy has a partition with B clusters, the next level has a partition with B^2 clusters, and so on. In practice, it is a data-dependent issue as to whether any of these clusters are empty.

11 Sparse Encoding of the Hierarchical Clustering

The regular B-way tree may be refined, using linear computational time processing, in order to regular $B-1$-way tree, $B-2$-way tree, and continuing. The aim, in this stepwise refinement, to exploit sparsity in the data.

It is interesting to find empirically that p-adic encoding, p prime, may best represent, i.e. approximate, the given hierarchical tree, [16].

12 Optimal Dimensionality Reduction of the Dual Cloud Spaces

Just to begin, consider our point cloud of observations in the attribute space, and, equivalently, our attribute point cloud in the observations space. Statistically optimal reduction of dimensionality, in Principal Component Analysis (PCA) for example, linearly transforms the point cloud exploiting those attributes that best preserve the variance of the point cloud. Such eigenreduction is of cubic computational complexity. Due to the dual space relationship, the determining of eigenvectors providing the latent, principal axes or factors, and the associated eigenvalues providing the variance explained by those axes, is carried out in either of our dual spaces. That is, carried out in either of the point clouds.

13 Clustering of a Point Cloud Inferred from Its Dual Space

Given the inherent relationships between the dual spaces, clustering of one cloud can be used to infer clusters in the dual space. Thus, we might consider clusters of attributes leading to clusters of their associated observations.

This can be considered as the basis for block mode clustering, see [11]. For statistical inference of clusters in a dual space, see also the `FactoMiner` package in R.

14 Exploiting Dual Space Relationships for Either $|I| \ll |J|$ or Vice Versa

Previous sections have considered how a seriation can be constructed whenever data piling arises, through massification of data. Informally, this might be expressed as follows. Data piling is compactification, or becoming condensed. The new origin of this mapped, or embedded, data is located at the centre of this piled or compactified data. The rescaling that is used will ensure that there is a point norm-based ordering.

Subject to setting up the data, as described in section 6, there will be this "condensation" of the data clouds resulting in: $f_{IJ} = f_I f_J$. When one of these clouds becomes very much compactified, that data piling can be viewed as approximating a single massive point. Geometrically this can be considered as becoming the product of a scalar and a vector. This therefore points to the relevance of one of the marginal distributions as defining the unidimensional mapping that is sought, to be the seriation to be used for the clustering.

15 Summarizing These Approaches to Dual Cloud Embedding, to Derive a Seriation of the Observations

Previous sections can be empirically verified. Key elements include: relative very large dimensionality; the normalization and rescaling that are used on the data prior to mapping to a Euclidean metric endowed space; normalization or data recoding that is used in the mapping.

One further aspect of importance is are distributional characteristics of the data, at all stages of the processing. This has a practical aspect also, in that applicability to massive data volumes gives rise to the need to appropriately encode one's data. Part and parcel of data encoding is the resolution scale of the data. Our use of hierarchical clustering is strongly motivated by this requirement for practical adaptability in applications.

It may be therefore necessary in practical applications to check on, and to monitor, data distributional properties, in order to benefit if distributional configurations are particularly simple, or if distributional configurations have problematic features.

16 From Seriation to Hierarchical Clustering

Hierarchical clustering through direct reading, for a given resolution level or for a given partition, can be viewed as quantization of the distribution of mapped observations onto the seriation structure. Traditional clustering, model-based or k-means, involves optimal quantization with non-fixed thresholds. Then in non-uniform quantization, with each label we associate a codebook entry, or associated cluster mean. See [13] for extensive discussion.

The regular hierarchical tree that results from the Baire metric and ultrametric provides quantization that has non-fixed thresholds relative to the clusters that are formed.

The Baire hierarachy or tree that results is an effective and efficient clustering method; it is adaptive; it is highly adaptable and adaptive for massive data sets; unlike traditional hierarchical clustering which may have partitions derived from it, in effect the Baire hierarchy or tree is constructed from the succession, in a top-down manner, of partitions.

As noted following section 3, precision relative to measured interpoint distances is not our objective. In a very massive data context, such precision is not a primary interest. Instead, let the following informal statement be considered. In astronomy, that deals with observational data from the cosmos, it is established and usual practice to see objects and clusters of objects that are deemed to be candidates, that is, candidates for selective, detailed or close-up further analytics.

Further Examples of Clustering through Quantization.

We conclude with a further commentary on quantization as an approach to clustering. In [17], there is the following.

Using our approach on the Fisher iris data, [6], 150 flowers crossed by petal and sepal width and breadth, provides the following outcome. We determine row sums, of the initial 150×4 data matrix, and the mean random projection of projections on 100 uniformly generated axes. From our previous results, we know that these are very highly correlated. We construct hierarchical clusterings on (i) the original 150×4 data matrix, (ii) the mean random projection, and (iii) the row sums. The cophenetic correlation coefficient is determined. (This uses ultrametric distances derived from the hierarchical tree, or dendrogram.) We find the cophenetic correlation of the hierarchies constructed on the row sums, and on the mean random projection, to be equal to 1 (as anticipated). Then between the hierarchy constructed on the 150×4 data matrix, and the mean random projection, the cophenetic correlation coefficient is 0.8798. For the given data and the row sums, it is 0.9885. The hierarchical clustering used was the average method; and other methods, including single link, provided very similar results. The distance used, as input to the hierarchical agglomerative clustering, was the square root of the squared Euclidean distance. Other alternatives were looked at, from the point of view of the distance used, and from the point of view of the agglomerative hierarchical clustering criterion.

We also looked at uniformly distributed, on [0,1], data of dimensions 2500×12. The correlation between row sums and mean of 100 random projections was 0.99. However, for the correlation between the hierarchical clustering on the original data, and the mean random projection, this correlation was 0.58. The correlation with the row sums was 0.578. The performance on this randomly generated data is seen to be not as good as that on the real valued, Fisher data. For data which is not strongly clustered, quantization is relevant. In the k-means clustering (partitioning) context, see e.g. [12]. Descriptively expressed, in quantization, in addition to cluster compactness, approximating identical cluster sizes is an objective.

gene-id	GSM177577	GSM177578	GSM177579	GSM177580	GSM177581	GSM177582
U48705.01	4091.2	3683.3	3117.9	3775.1	3510.3	3461.4
M87338.02	762.7	666.9	581.2	623.5	822.9	735.9
X51757.03	113.3	112.8	90.5	128	120.7	93.1

Table 1: A sample of the data. Genes crossed by samples.

17 Practical Case Studies of Mapping Seriation to Hierarchy

In [18], it is shown on Fisher's iris data how hierarchical clustering using a seriation that is the aggregate or consensus random projection, and the row mass distribution, is well and truly associated (high cophenetic correlation, that is, correlation of the ultrametric distances) with an agglomerative hierarchical clustering method.

18 Case Study, Towards Both Partition-Based Clustering and Hierarchical Clustering with Linear Computational Complexity

We draw on the wide-ranging vantage points offered by the methodologies for carrying out clustering, and related analytics processing. The latter may possibly include: dimensionality reduction, orthonormal factor or principal axis space mapping, feature or attribute selection, and so on. Effectively we are drawing on what has been overviewed in previous parts of this paper.

18.1 Step 1: Determining the Seriation

Gene expression data are used, where rows contain genes and columns contain samples, example in Table 1.

The genes dataset used is a table of dimensions 61359×16. The data, x_{ij} values for $i \in I, |I| = 61359, j \in J, |J| = 16$, are extremely exponentially distributed in their values; the row sums, x_I such that $x_i = \sum_{j \in J} x_{ij}$, are also extremely exponentially distributed; and the column sums, x_J such that $x_J = \sum_{i \in I} x_{ij}$, is Gaussian. For the latter, the Shapiro-Wilk normality test gives a p-value of 0.9145. Just to illustrate values of $x_{IJ} = x_{ij} \forall i \in I, j \in J$ we ave the maximum, minimum, mean and median as follows: 676496, 0.199789, 1130.097, 130.5835.

Next checking the mean of random projections, with uniformly distributed axis coordinate values, we find very high correlation with the row masses ($f_I = x_I / \sum_{i,j} x_{ij}$), equivalent for correlation to the row sums, x_I. Note the ordinate scale, that starts, with just one random projection, at a correlation of almost 0.999.

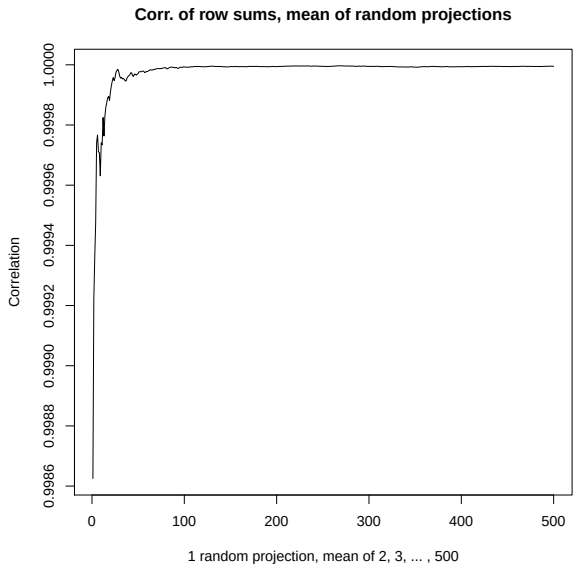

Figure 1: Correlation between the row sums, and the mean random projection. The random projections are the 1st, then the mean of the 1st and 2nd, the mean of the 1st, 2nd, and 3rd; and so on, up to the mean of the 1st, 2nd, 3rd, 4th, ..., 500th random projections.

This first finding indicates how very consistent the scaling is, when taking the seriation of genes, that is provided by any of the following: (i) distributional masses of the genes (notationally here, f_I).

18.2 Step 2: Transforming or Re-Encoding the Seriation

Arising out of the previous subsection, exemplifying the simplification property of cloud one of the clouds $\mathcal{N}(I)$, $\mathcal{N}(J)$ stated in 6, we proceed to examining $f_I = \{f_i | i \in I\}$ where $f_i = \sum_{j \in J} x_{ij} / \sum_{i \in I, j \in J} x_{ij}$.

The seriation is to be directly used for deriving a partition of clusters or a hierarchy of clusters. This is to be carried out with computational time that is linear or $O(n)$ for n observations, using earlier notation, $n = |I|$.

The distribution of the seriation values is found to be highly exponentially decreasing. A direct, density-based, determining of clusters could lead to just one single cluster with all observations as members of it, but this would be futile in practice. Hence we would like to re-encode our seriation values to carry out quantization-based determining of clusters. Cf. section 16. We could sort our seriation values and then read the values, for a top-level partition with 10 clusters, as a series of 10%, 20%, etc. quantiles. However, sorting the n-length set of seriation values requires $O(n \log n)$ computational complexity. We see rather to bound our processing to be $O(n)$.

Due to exponentially distributed values, we first take the log of these values. Since this is Gaussian distributed, we standardize it to zero mean and unit standard deviation.

Next we uniformize the standard Gaussian distributed values, i.e. we convert to a uniform distribution, using the complementary error function. Complementary error function is one minus the error function; the error function is the probability of a value being within the range $0, x/\sigma\sqrt{2}$ The error function is twice the integral of a normalized Gaussian function in this interval. Cf. [5].

The effectiveness of these re-encodings is displayed in the histograms in Figure 2. Our aim is to have a reasonable strategy that is computationally of linear complexity, for taking the highly exponentially distributed values into a near uniformly distributed set of values. By a reasonable strategy, we mean an approach which is generalizable, which is assessed at least visually using histograms, and also which can be further considered or indeed, subject to alternative generalizable and computationally scalable approaches.

18.3 Step 3: Reading off Partitions at Successive Hierarchical Levels

We directly read off the cluster members by taking the first, viz. top, partition as consisting of 10 clusters. Our uniformized data values are in the closed/open interval, $[0, 1)$. We find the following cluster cardinalities, respectively for the clusters that are labelled 0, 1, 2, 3, 4, 5, 6, 7, 8, 9: 2345, 8445, 10347, 9316, 6948, 5122, 3999, 3588, 3926, 7323.

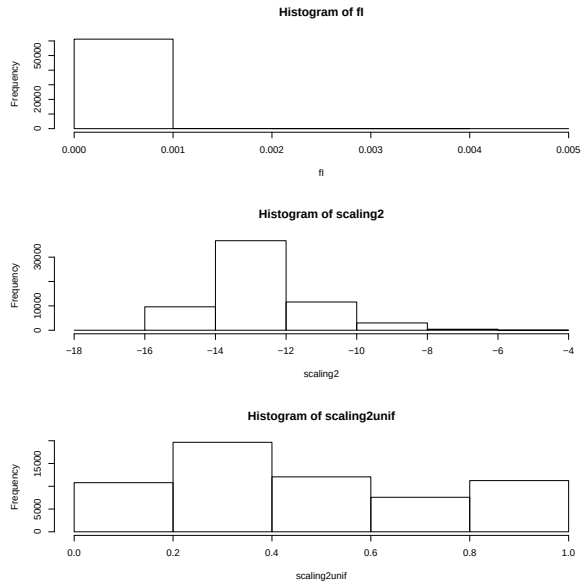

Figure 2: Histograms of marginal distribution; of standard normal distributed re-encoding; and of uniformized re-encoding.

Equally straightforwardly we can read off the next partition in a regular 10-way hierarchy or tree, that define the clusters with labels 0, 1, 2, ..., 98, 99. Following the top partition with 10 clusters, that second level partition with 100 clusters, then the next, third level, partition would have 1000 clusters, and so on.

In the Appendix, the R code used is listed.

18.4 Cluster Properties of Top Level Partition

In order to validate the clustering, here the properties of the top level partition, with 10 clusters, are examined. The distribution of the given data, as noted above, is exponentially decreasing. We determine the mean vectors of the clusters in the given 16-dimensional attribute space. Next, all pairwise distances are obtained. These are Euclidean distances. For each cluster, the nearest neighbour cluster is determined. Because of the distributional properties of our data, we find that cluster 0 has as its nearest neighbour, cluster 1; cluster 1 has as its nearest neighbour, cluster 2; and so on, in succession. These nearest neighbour clusters, and the nearest neighbour distances, are listed in Table 2.

To summarize how compact the clusters are, for each cluster we determine the variances in each dimension, aggregate (sum) these variances to have the overall cluster variance, and then list the overall cluster standard deviation, just multiplied by 3, in Table 2. The 3σ, i.e. 3 standard deviation, measure used

Clusters	3 * Std.dev	NN-cluster	NN-dist.
0	49.92530	1	94.48774
1	85.11277	2	125.75401
2	124.10223	3	160.40965
3	130.90670	4	213.86643
4	161.75631	5	296.95453
5	205.36007	6	445.92892
6	180.60822	7	767.96664
7	275.07759	8	1586.48549
8	367.02603	9	27694.34381
9	1152.21181		

Table 2: The top or first level partition, with 10 clusters. Column 2 lists 3× standard deviation of the cluster. Column 3 lists the nearest neighbour cluster of each cluster. Finally column 4 lists that nearest neighbour distance.

is just so as to have comparability with what would be relevant for a Gaussian distributed cluster.

What Table 2 demonstrates very well, is that cluster compactness, measured by the 3σ column, is far less in value than the cluster-to-cluster nearest neighbour distance. Therefore this leads to the conclusion that the cluster properties, in this top level partition into 10 clusters, are excellent.

19 Conclusion

The clustering approach used here is quite simple to implement, and its methodological basis has been straightforwardly and briefly described in previous sections. In section 16 in particular, there is the brief description of the methodological foundations for this approach to hierarchical clustering and determination of one or more partitions.

While it may be noted that instead of a compactness clustering criterion, we are using quantization as the main basis for the clustering. For candidate object selection, and so on, as is common in observational science, then this can be claimed to be quite adequate, as a methodology. If there are physical laws, however, at issue, then these may be used, e.g. through feature selection, or through statistical modelling. For the latter, we have noted throughout this work just what distributional properties were found, or were verified, to hold.

While this work has considered the data properties, it is to be considered therefore that this particular analytics processing pipeline, that has been the focus of this work, is associated with data of a given class or family of distributional and related properties.

A major justification for our extensive description of methodological underpinnings is to allow for alternative, but related, approaches to be designed and implemented when there are somewhat different ultimate objectives. This might

include the incorporation of supervised learning phases in the analytics chain or pipeline.

References

[1] J.P. Benzécri, "L'approximation stochastique en analyse des correspondances", *Les Cahiers de l'Analyse des Données*, 7(4), 387–394, 1982.

[2] P. Contreras, *Search and Retrieval in Massive Data Collections*, PhD, Royal Holloway University of London, 2010.

[3] P. Contreras and F. Murtagh, "Fast, linear time hierarchical clustering using the Baire metric", *Journal of Classification*, 29, 118–143, 2012.

[4] F. Critchley and W. Heiser, "Hierarchical trees can be perfectly scaled in one dimension", *Journal of Classification*, 5, 5–20, 1988.

[5] J.R. Culham, "Error and complementary error functinos", lecture notes, http://www.mhtlab.uwaterloo.ca/courses/me755/web_chap2.pdf

[6] R.A. Fisher, "The use of multiple measurements in taxonomic problems", *The Annals of Eugenics*, 7, 179–188, 1936.

[7] P. Hall, J.S. Marron and A. Neeman, "Geometric representation of high dimension, low sample size data", *Journal of the Royal Statistical Society Series B*, 67, 427–444, 2005.

[8] W.B. Johnson and J. Lindenstrauss, "Extensions of Lipschitz maps into a Hilbert space", *Conference in Modern Analysis and Probabilities. 26. Contemporary Mathematics*, Providence RI: American Mathematical Society, pp. 189–206, 1984.

[9] S. Kaski, "Dimensionality reduction by random mapping: fast similarity computation for clustering", *Proceedings of The 1998 IEEE International Joint Conference on Neural Networks*, pp. 413–418, 1998.

[10] P. Li, T. Hastie and K. Church, "Very sparse random projections", *KDD 2006: Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Volume 1, New York: ACM, 287–296, 2006.

[11] I. Liiv, "Seriation and matrix reordering methods: An historical overview", *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 3 (2), 70–91, 2010.

[12] S.P. Lloyd, "Least-squares quantization in PCM", *IEEE Transactions on Information Theory*, IT-28, 129–137, 1982.

[13] F. Murtagh and J.L. Starck, "Quantization from Bayes factors with application to multilevel thresholding", *Pattern Recognition Letters*, 24, 2001–2007, 2003.

[14] F. Murtagh, G. Downs and P. Contreras, "Hierarchical clustering of massive, high dimensional data sets by exploiting ultrametric embedding", *SIAM Journal on Scientific Computing*, 30, 707–730, 2008.

[15] F. Murtagh and P. Contreras, "Random projection towards the Baire metric for high dimensional clustering", in A. Gammerman et al., *Proc. Symposium on Learning and Data Sciences*, Springer LNAI Vol. 9047, pp. 424–431, 2015.

[16] F. Murtagh, "Sparse p-adic data coding for computationally efficient and effective big data analytics", *p-Adic Numbers, Ultrametric Analysis and Applications*, 8(3), 236-247, 2016

[17] F. Murtagh and P. Contreras, "Clustering through high dimensional data scaling: Applications and implementations", *Archives of Data Science, Series A*, 2(1), 1-16, 2017.

[18] F. Murtagh and P. Contreras, "Direct reading algorithm for hierachical clustering", *Electronic Notes in Discrete Mathematics*, 6, 37–42, 2016.

[19] Y. Terada, "Clustering for high-dimension, low-sample size data using distance vectors", 16 pp., 2013, http://arxiv.org/abs/1312.3386

Appendix

In the following R code, we show dimensions and other data properties, such as maximum and minimum values, and mean and median, in order to both indicate data values, and to support reproducibility.

```
> x <- read.table("dataset-1.tsv", header = TRUE, row.names=1); dim(x) # 61359 x 16
> x2 <- as.matrix(x)  # As read, x is of list data type.
> max(x2); min(x2); mean(x2); median(x2)  # 676496 0.199789 1130.097 130.5835

# Determine the marginal distribution; then re-encode from its very exponential
# distribution.
> fI <- apply(x2, 1, sum)/sum(x2)
> min(fI); max(fI)  #  1.049568e-07 0.004014821
# Exp to normal, then normal to uniform:
> scaling2 <- log(fI)
> scaling2norm01 <- (scaling2 - mean(scaling2))/sd(scaling2)
# Next use complementary error function.  Cf. help(Normal)
> erfc <- function(x) 2 * pnorm(x * sqrt(2), lower = FALSE)
> scaling2unif <- 0.5*erfc(- scaling2norm01/sqrt(2))
```

```
> max(scaling2unif); min(scaling2unif)  #  0.9999998  0.01001248

# Top level partition with 10 clusters labelled 0, 1, ... , 9.
> scaling2class10 <- trunc(scaling2unif * 10)
# Next level partition with 100 clusters labelled 0, 1, ... , 99.
# Note that we find the cluster with label 0 to be empty.
> scaling2class100 <- trunc(scaling2unif * 100)
> scaling2class1000 <- trunc(scaling2unif * 1000) # 3rd level partition.

# Cluster cardinalities:
> table(scaling2class10)  # Top level partition, 10 clusters.
scaling2class10
    0     1     2     3     4     5     6     7     8     9
 2345  8445 10347  9316  6948  5122  3999  3588  3926  7323
> table(scaling2class100)  # Second level partition, <= 100 clusters.
> table(scaling2class1000)  # Third level partition, <= 1000 clusters.
> length(table(scaling2class100))  # 99 (non-empty clusters)
> length(table(scaling2class1000))  # 989 (non-empty clusters)
```